What is clained is:

1. A method of smoothing fundamental frequency discontinuities at boundaries of concatenated speech segments, each speech segment characterized by a segment fundamental frequency contour and including two or more frames, comprising:

determining, for each speech segment, a beginning fundamental frequency value and an ending fundamental frequency value;

adjusting the fundamental frequency contour of each of the speech segments according to a predetermined function calculated for each particular speech segment, wherein parameters characterizing each predetermined function are selected according to the beginning fundamental frequency value and the ending fundamental frequency value of the corresponding speech segment.

2. A method according to claim 1, wherein the predetermined function adjusts a slope associated with the speech segment.

3. A method according to claim 1, wherein the predetermined function adjusts an offset associated with the speech segment.

4. A method according to claim 1, wherein the predetermined function includes a linear function.

5. A method according to claim 1, wherein the predetermined function calculated for each particular speech segment is dependent upon a length associated with the speech segment, such that the predetermined function adjusts longer segments more than shorter segments.

6. A method according to claim 1, further including determining, for each speech segment one or more parameters selected from: (i) a total duration of the segment; (ii) a total duration of all voiced regions of the segment; (iii) a average value of the fundamental frequency contour over all voiced regions of the segment; (iv) a median value of the fundamental frequency contour over all voiced regions of the segment; and (v) a standard deviation of the fundamental frequency contour over the whole segment.

7. A method according to claim 6, further including setting the determined median value of the fundamental frequency contour over all voiced regions of the segment to the average value of the fundamental frequency contour over all voiced regions of the segment if a number of fundamental frequency samples in the speech segment is less than a predetermined value.

8. A method according to claim 1, further including examining a predetermined number of frames from a beginning point of each speech segment, and setting the beginning fundamental frequency value to a fundamental frequency value of the first frame if all fundamental frequency values of the predetermined number of frames from the beginning point of the speech segment are within a predetermined range.

9. A method according claim 1, further including examining a predetermined number of frames from an ending point of each speech segment, and setting the ending fundamental frequency value to a fundamental frequency value of the last frame if all fundamental frequency values of the predetermined number of frames from the ending point of the speech segment are within a predetermined range.

10. A method according to claim 1, further including setting the beginning fundamental frequency and the ending fundamental frequency of unvoiced speech segments to a value substantially equal to a median value of the fundamental frequency contour over all voiced regions of a preceding voiced segment.

11. A method according to claim 1, further including calculating, for each pair of adjacent speech segments n and n+1 one or more of: (i) a first ratio of the n^{th} ending fundamental frequency value to the $n+1^{th}$ beginning fundamental frequency value; and (ii) a second ratio being the inverse of the first ratio; and adjusting the n^{th} ending fundamental frequency value and the $n+1^{th}$ beginning fundamental frequency value only if the first ratio and/or the second ratio are less than a predetermined ratio threshold.

12. A method according to claim 1, further including calculating the function for each individual speech segment according to a coupled spring model.

13. A method according to claim 12, further including implementing the coupled spring model such that a first spring component couples the beginning fundamental frequency value to an anchor component, a second spring component couples the ending fundamental frequency value to the anchor component, and a third spring component couples the beginning fundamental frequency value to the ending fundamental frequency value.

14. A method according to claim 13, further including associating a spring constant with the first spring and the second spring such that the spring constant is proportional to a duration of voicing in the associated speech segment.

15. A method according to claim 13, further including associating a spring constant with the third spring such that the third spring models a non-linear restoring force that resists a change in slope of the segment fundamental frequency contour.

16. A method according to claim 12, further including forming a set of simultaneous equations corresponding to the coupled spring models associated with all of the concatenated speech segments, and solving the set of simultaneous equations to produce the parameters characterizing each linear function associated with one of the speech segments.

17. A method according to claim 16, further including solving the set of simultaneous equations through an iterative algorithm based on Newton's method of finding zeros of a function.

18. A system for smoothing fundamental frequency discontinuities at boundaries of concatenated speech segments, each speech segment characterized by a segment fundamental frequency contour and including two or more frames, comprising:

a unit characterization processor for receiving the speech segments and characterizing each segment with respect to a beginning fundamental frequency and an ending fundamental frequency;

a fundamental frequency adjustment processor for receiving the speech segments, the beginning fundamental frequency and ending fundamental frequency, and for adjusting the fundamental frequency contour of each of the speech segments according to a predetermined function calculated for each particular speech segment, wherein parameters characterizing each predetermined function are selected according to the beginning fundamental frequency value and the ending fundamental frequency value of the corresponding speech segment.

19. A system according to claim 18, wherein the predetermined function adjusts a slope associated with the speech segment.

20. A system according to claim 18, wherein the predetermined function adjusts an offset associated with the speech segment.

21. A system according to claim 18, wherein the predetermined function includes a linear function.

22. A system according to claim 18, wherein the predetermined function calculated for each particular speech segment is dependent upon a length associated with the speech segment, such that the predetermined function adjusts longer segments more than shorter segments.

23. A system according to claim 18, wherein the unit characterization processor determines, for each speech segment one or more of: (i) a total duration of the segment; (ii) a total duration of all voiced regions of the segment; (iii) an average value of the fundamental frequency contour over all voiced regions of the segment; (iv) a median value of the fundamental frequency contour over all voiced regions of the segment; and (v) a standard deviation of the fundamental frequency contour over the whole segment.

24. A system according to claim 23, wherein the unit characterization processor sets the determined median value of the fundamental frequency contour over all voiced regions of the segment to the average value of the fundamental frequency contour over all voiced regions of the segment if a number of fundamental frequency samples in the speech segment is less than a predetermined value.

25. A system according to claim 18, wherein the unit characterization processor examines a predetermined number of frames from a beginning point of each speech segment, and sets the beginning fundamental frequency value to a fundamental frequency value of the first frame if all fundamental frequency values of the predetermined number of frames from the beginning point of the speech segment are within a predetermined range.

26. A system according to claim 18, wherein the unit characterization processor examines a predetermined number of frames from a ending point of each speech segment, and sets the ending fundamental frequency value to a fundamental frequency value of the last frame if all fundamental frequency values of the predetermined number of frames from the ending point of the speech segment are within a predetermined range.

27. A system according to claim 18, wherein the unit characterization processor sets the beginning fundamental frequency and the ending fundamental frequency of unvoiced speech segments to a value substantially equal to a median value of the fundamental frequency contour over all voiced regions of a preceding voiced segment.

28. A system according to claim 18, wherein the unit characterization processor calculates, for each pair of adjacent speech segments n and n+1 one or more of: (i) a first ratio of the n^{th} ending fundamental frequency value to the $n+1^{th}$ beginning fundamental frequency value; and (ii) a second ratio being the inverse of the first ratio, and adjusts the n^{th} ending fundamental frequency value and the $n+1^{th}$ beginning fundamental frequency value only if the first ratio and/or the second ratio are less than a predetermined ratio threshold.

29. A system according to claim 18, wherein the fundamental frequency adjustment processor calculates the linear function for each individual speech segment according to a coupled spring model.

30. A system according to claim 29, wherein the fundamental frequency adjustment processor implements the coupled spring model such that a first spring component couples the beginning fundamental frequency value to an anchor component, a second spring component couples the ending fundamental frequency value to the anchor component, and a third spring component couples the beginning fundamental frequency value to the ending fundamental frequency value.

31. A system according to claim 30, wherein the fundamental frequency adjustment processor associates a spring constant with the first spring and the second spring such that the spring constant is proportional to a duration of voicing in the associated speech segment.

32. A system according to claim 30, wherein the fundamental frequency adjustment processor associates a spring constant with the third spring such that the third spring models a non-linear restoring force that resists a change in slope of the segment fundamental frequency contour.

33. A system according to claim 29, wherein the fundamental frequency adjustment processor forms a set of simultaneous equations corresponding to the coupled spring models associated with all of the concatenated speech segments, and solves the set of simultaneous

equations to produce the parameters characterizing each linear function associated with one of the speech segments.

34. A system according to claim 33, wherein the fundamental frequency adjustment processor solves the set of simultaneous equations through an iterative algorithm based on Newton's method of finding zeros of a function.

36 A method of smoothing fundamental frequency discontinuities at boundaries of concatenated speech segments, each speech segment characterized by a segment fundamental frequency contour and including two or more frames, comprising:

adjusting the fundamental frequency contour of each speech segment according to a predetermined function calculated for each particular speech segment, wherein the predetermined function is dependent upon a length associated with the speech segment, such that the predetermined function adjusts longer segments more than shorter segments.

37. A system for smoothing fundamental frequency discontinuities at boundaries of concatenated speech segments, each speech segment characterized by a segment fundamental frequency contour and including two or more frames, comprising:

a fundamental frequency adjustment processor for adjusting the fundamental frequency contour of each speech segment according to a predetermined function calculated for each particular speech segment, wherein the predetermined function is dependent upon a length associated with the speech segment, such that the predetermined function adjusts longer segments more than shorter segments.